Shepherd's Dilemma

Selected Paper prepared for presentation at the

American Agricultural Economics Association Annual Meeting

Providence, Rhode Island, July 24-27, 2005

Glenn Sheriff[1] and Daniel Osgood[2]

May 15, 2005

[1] Columbia University SIPA
[2] Columbia University IRI

Abstract

Recent outbreaks of Rift Valley Fever in sheep have led to boycotts of African livestock by Middle Eastern importers. To normalize trade, attempts have been made to apply new livestock forecasting and monitoring technologies. In this process, producers have exhibited a resistance in revealing livestock health information, a resistance that could jeopardize the information system and lead to further boycotts. We investigate the incentives governing this problem and model the most fundamental contract issues, those concerning reputation and credibility. Equilibrium contracts require that the buyer compensate the producer for private information to address the shepherd's dilemma of concealing livestock information (and facing continued boycotts) or revealing the information and being blacklisted.

1 Introduction

The livelihoods of many nomadic pastoralists in the Greater Horn of Africa depend on the export of lambs to the Middle East for use in religious ceremonies. Recent outbreaks of Rift Valley Fever (RVF), a relatively new livestock disease that can be transmitted to humans through ceremonial contact with infected sheep, have led to boycotts by importing countries. RVF is rare, at most impacting a small fraction of the exported sheep. It is, however, highly feared and deadly when communicated to humans. This market collapse has been very costly to all concerned: consumers must resort to Australian lambs of inferior quality and higher price, while producers have seen their livelihoods threatened (Consultative Meeting of Experts on Rift Valley Fever, 2001) .

Some market participants have enlisted international groups to develop an RVF forecasting and monitoring system in order to normalize trade and prevent future boycotts (Dilley, 2003). This system could reduce the likelihood of contaminated lambs entering the market and thereby eliminate the need for continent-wide quarantines. Two technological advances may make the forecasting/monitoring system feasible. First, recent advances in veterinary medicine may lower the cost of testing for RVF. Regardless of how inexpensive the monitoring procedure becomes, however, the cost of obtaining timely samples from sheep across the entire Greater Horn of Africa remains prohibitive.

Considerable potential exists for using new technologies in seasonal precipitation forecasts to predict RVF outbreaks (Indeje et al., 2004). Due to its reliance on the mosquito as a vector, the incidence of RVF is correlated with precipitation patterns. If successful, this technology could allow a monitoring system to target its sample collection and greatly reduce costs, perhaps making such as system economically feasible. In addition, the forecasts might provide information useful to producers in performing risk mitigating strategies, such as minimizing grazing in high-risk areas and immunizing targeted livestock in high-risk years.[1]

Cooperative participation of the producers is essential for any RVF management system. In spite of the potential benefits, however, producers have been reluctant to cooperate in the development of an RVF forecasting and monitoring system. The presence RVF in a flock is evident to a producer due to highly visible symptoms such as unusually large numbers of sheep abortions. These signs are unobservable to importers, however, due to the extreme geographic isolation of the producers. Once the lamb has been

[1] The RVF immunizations being considered are relatively expensive or cause damaging side effects to livestock. They must be applied weeks before livestock are exposed to the disease to be effective and immunity benefits may decay quickly over time(Consultative Meeting of Experts on Rift Valley Fever, 2001).

brought to market, RVF is costly to detect. For prediction, in order to successfully model the climate-disease interaction it is necessary to obtain local data concerning RVF episodes directly from the producers. In addition, if implemented, the forecasting/monitoring system would need the continued cooperation of livestock producers for symptoms reporting and access to animals for testing. Finally, because there is a potential for violence in many of the impacted regions, cooperation is especially worthwhile.

At its most basic level, this problem is reminiscent of Akerlof's (1970) lemons model. The pastoralists have private information regarding the health of their lambs. The likelihood of a lamb being sick is common knowledge. If this probability is high enough, the buyer is unwilling to offer a price high enough to induce owners of healthy lambs to sell and the market collapses.

Signaling by high quality sellers is often used as a means of overcoming the lemons problem. Why then are sellers so opposed to development of an RVF forecasting system? One reason is that individual producers fear that if they inform researchers of an RVF outbreak, they will be blacklisted. Those most vulnerable to the market impacts of RVF may be made more vulnerable by the international forecasting/monitoring effort.[2]

In this paper we present the most central elements to the problem, the fundamental issues of reputation and credibility that must be overcome for the market to function. For clarity and ease of generalization, we present the problem in the simplest form possible, focusing on the interaction between a single buyer and a population of sellers with private information.

2 The Model

The model builds upon the literature of optimal contracts with costly information verification pioneered by Townsend (1979). In a typical example, an uninformed principal, such as an investor, initiates a contract with an agent (entrepreneur) about to undertake a risky project. Only the entrepreneur knows the precise outcome of the project. The investor can obtain the agents' information, but only with a costly audit. The principal's task is to design a contract that stipulates the transfers to be paid and probability of audit depending on the agent's declared income. The optimal contract maximizes the net income of the principal.

We model the interaction between the buyer (principal) and the producers (agents) as a repeated game

[2](Pfaff et al., 1999; Osgood, 2005) have found evidence in other contexts that forecasting and monitoring information may impose hardships on specific subgroups of a population.

with asymmetric information. The basic forces driving the contractual features can be demonstrated in analysis of two players in two periods. In Townsend's initial model the interaction between the parties was not repeated, and the underlying motivation for the contract was one of risk sharing. In order to focus on the market for lambs rather than an insurance, we follow the work of Border and Sobel (1987) and assume that both the buyer and the pastoralist are risk neutral.[3]

In our context, the importer is the principal. Acting as a Stackelberg leader, she proposes a menu of contracts to the agents (pastoralists). Each agent chooses the contract that maximizes his expected income. We obtain the Bayesian-Nash equilibrium to the game by backwards induction. Knowing how the agents will respond to any given contract menu, the principal designs the menu that maximizes her expected income.

Each agent has one lamb. The status of the lamb (sick or healthy) determines an agent's type, $\theta \in \Theta \equiv \{\theta_s, \theta_h\}$, and is known only to the agent. Type is perfectly persistent across time periods. The reservation income to the agent of keeping, rather than selling, the lamb is a function of its type, $r(\theta)$, where $r(\theta_h) > r(\theta_s)$.

The value to the principal of obtaining a lamb is a function of type, $P(\theta)$, where $P(\theta_h) \geq r(\theta_h) \geq r(\theta_s) \geq P(\theta_s)$. Thus, ideally, the principal would like only to purchase lambs from healthy types. The principal's prior belief regarding the probability that an agent's lamb is sick is $\rho \in [0, 1]$. The principal can verify the type of any given lamb by conducting an "audit" with cost a.[4]

In the first period, each agent's set of possible actions consists of selecting a contract. We restrict attention to contracts that depend only upon an individual agent's actions. That is to say, an agent's payoff from making a declaration does not depend upon other agents' declarations. Under this restriction, the revelation principle ensures that there is no loss in generality by restricting attention to truthful direct revelation mechanisms, whereby each agent's action is redefined as making a declaration $\tilde{\theta}_1 \in \Theta$ regarding his type (Townsend, 1988).

The first period contracts consists of a probability of audit $G_1\left(\tilde{\theta}_1\right) \in [0, 1]$ that is a function of an agent's declaration, and a set of three possible transfers. If the agent is not audited, he receives transfer $t_1\left(\tilde{\theta}_1\right)$. If an audited agent is telling the truth, he receives $\bar{t}_1\left(\tilde{\theta}_1\right)$ rather than $t_1\left(\tilde{\theta}_1\right)$. If an agent is caught lying, he receives $\underline{t}_1\left(\tilde{\theta}_1\right)$ rather than $t_1\left(\tilde{\theta}_1\right)$. The set of first period contracts terms consists of two possibilities

[3] Mookherjee and Png (1989) and Wang (2005) respectively consider problems of costly state verification with risk averse agents in a single period and dynamic settings.

[4] For the more general case, the average cost of auditing could depend on the accuracy of a forecast as well as the cost of a medical test, since a skillful forecast could allow one to target medical sampling.

presented to each agent: $G_1\left(\tilde{\theta}_1\right), t_1\left(\tilde{\theta}_1\right), \bar{t}_1\left(\tilde{\theta}_1\right), \underline{t}_1\left(\tilde{\theta}_1\right)$, for $\tilde{\theta}_1 \in \Theta$.

The agent action space in the second period is also simply declaration of type $\tilde{\theta}_2 \in \Theta$. The set of contracts is more complicated, however, since the audit probabilities and transfers are contingent not only upon $\tilde{\theta}_2$, but also information obtained in the first period. Specifically, by the revelation principle, an agent's action in the first period reveals his true type, regardless of whether he is audited. The set of four possible second period contracts is: $G_2\left(\tilde{\theta}_1, \tilde{\theta}_2\right), t_2\left(\tilde{\theta}_1, \tilde{\theta}_2\right), \bar{t}_2\left(\tilde{\theta}_1, \tilde{\theta}_2\right), \underline{t}_2\left(\tilde{\theta}_1, \tilde{\theta}_2\right)$, for $\left(\tilde{\theta}_1, \tilde{\theta}_2\right) \in \Theta^2$.

We assume that the agents have limited liability. As a result, all transfers are bounded from below by zero. In addition, participation in the contracts is voluntary ex post. Therefore no agent can receive less than his reservation income by choosing a contract.

Letting $\delta > 0$ be the discount rate, the present value of the principal's expected welfare from a menu of truthful contracts is:

$$\begin{aligned} W = {} & [1-\rho]\left\{P(\theta_h) - [G_1(\theta_h)[\bar{t}_1(\theta_h) + a] + [1 - G_1(\theta_h)]\, t_1(\theta_h)]\right. \\ & \left. +\delta\left[P(\theta_h) - [G_2(\theta_h,\theta_h)[\bar{t}_2(\theta_h,\theta_h) + a] + [1 - G_2(\theta_h,\theta_h)]\, t_2(\theta_h,\theta_h)]\right]\right\} \\ & +\rho\left\{P(\theta_s) - [G_1(\theta_s)[\bar{t}_1(\theta_s) + a] + [1 - G_1(\theta_s)]\, t_1(\theta_s)]\right. \\ & \left. +\delta\left[P(\theta_s) - [G_2(\theta_s,\theta_s)[\bar{t}_2(\theta_s,\theta_s) + a] + [1 - G_2(\theta_s,\theta_s)]\, t_2(\theta_s,\theta_s)]\right]\right\} \end{aligned} \tag{1}$$

In a repeated game such as this, the agent's reputation is critical. In particular, the way that the principal uses the information revealed during the first period drives the equilibrium outcome. The ability of the principal to credibly commit to not using information obtained in the first period can lead to different equilibria.

If the principal cannot commit to ignoring information obtained in the first period, then the set of credible second period contracts is reduced. If an agent's type is known, it is not reasonable for the agent to believe that the principal will incur the cost of an audit in the second period, pay more than the healthy agent's reservation income, or offer a contract to a sick agent. Consequently, without the ability to commit we have at equilibrium $G_2(\theta,\theta) = 0$, for $\theta \in \Theta$, $t_2(\theta_h,\theta_h) = r(\theta_h)$, and $t_2(\theta_s,\theta_s) = 0$.

Since the set of credible second period contracts is reduced, the principal cannot be made worse off by being able to commit to not using information obtained in the first period. Thus, the principal may benefit if she can commit to "forgetting" information from the first period. To address the impact of commitment and reputation, we first determine the optimal contract menu for the principal to offer assuming she is able to commit to ignoring first-period information. We then consider the case where she is unable to make this commitment.

2.1 Perfect commitment

It is a standard result that when intertemporal commitment is permitted and agent types are perfectly persistent across time, the Bayesian-Nash equilibrium of a multi-period game is simply the repeated equilibrium of the single period game (Baron and Besanko, 1984). For the single-period game, the principal chooses the contract terms that maximize expected current welfare:

$$\begin{aligned} W_1 = {} & [1-\rho]\left\{P(\theta_h) - [G_1(\theta_h)[\bar{t}_1(\theta_h)+a] + [1-G_1(\theta_h)]\,t_1(\theta_h)]\right\} \\ & + \rho\left\{P(\theta_s) - [G_1(\theta_s)[\bar{t}_1(\theta_s)+a] + [1-G_1(\theta_s)]\,t_1(\theta_s)]\right\} \end{aligned} \tag{2}$$

When confronted with a menu of contracts, an agent chooses the one that maximizes his expected income.[5] Let $I_1\left(\tilde{\theta},\theta\right)$ denote the expected income of an agent of type θ choosing the contract intended for an agent of type $\tilde{\theta}$. Since participation is voluntary, if a transfer is below an agent's reservation income he will not accept the transfer and keep his sheep (and his reservation income). Consequently, we have

$$I_1(\theta,\theta) = G_1(\theta)\max\{\bar{t}_1(\theta), r(\theta)\} + [1-G_1(\theta)]\max\{t_1(\theta), r(\theta)\}, \tag{3}$$
$$\text{for } \theta\in\Theta,$$

and

$$I_1\left(\tilde{\theta},\theta\right) = G_1\left(\tilde{\theta}\right)\max\left\{\underline{t}_1\left(\tilde{\theta}\right), r(\theta)\right\} + \left[1-G_1\left(\tilde{\theta}\right)\right]\max\left\{t_1\left(\tilde{\theta}\right), r(\theta)\right\}, \tag{4}$$
$$\text{for } \theta,\tilde{\theta}\in\Theta, \theta\neq\tilde{\theta}.$$

For the direct revelation mechanism to be truthful, each agent must voluntarily to choose his own contract. That is to say, the following incentive compatibility (IC) constraint must be satisfied for each type: $I_1(\theta,\theta)\geq$

[5] We assume that when indifferent between the contract intended for it by the principal and another contract, the agent chooses the intended contract.

$I_1\left(\tilde{\theta},\theta\right)$.

Note from the definition of principal welfare in Eq. (2), that all transfers are costly. Since reducing the penalty transfers $\underline{t}(\theta)$ increases principal welfare and weakens IC constraints, at the optimum these transfers are as low as possible:

$$\underline{t}_1(\theta) = 0 \text{ for all } \theta. \tag{5}$$

In this setting there are two classes of optimal contract menus. In one class, no auditing occurs and each type receives a payment equal to the healthy type's reservation income. In the other, healthy declarations are always audited and sick declarations are never audited. In this case, the principal only purchases from the healthy type, and each type receives his reservation income. Which of these two menus is optimal depends on the values of the reservation incomes $r(\theta_h)$ and $r(\theta_s)$, the probability of a sheep being sick ρ, and the cost of audit a.

This characterization can be proved in two stages. First, we show that it is never optimal for a sick declaration to be audited. Next, we show that it is either optimal to audit a healthy declaration with probability zero or one, depending on the exogenous variables.

Note that for any contract in which $G_1(\theta_s) > 0$, it is optimal to set the unaudited transfer $t_1(\theta_s)$ equal to zero. To see why, note that if $G_1(\theta_s) > 0$ and $t_1(\theta_s) > 0$, then $t_1(\theta_s)$ can be slightly reduced without changing the value of $I(\theta_s,\theta_s)$ by slightly increasing $\bar{t}_1(\theta_s)$ if necessary. Such a change would not reduce the principal's expected welfare but could weaken the IC constraint by reducing $I_1(\theta_s,\theta_h)$.

Suppose that it were in fact optimal to audit a sick declaration with positive probability. In that case, using Eq. (5)

$$I_1(\theta_s,\theta_h) = r(\theta_h). \tag{6}$$

Hence, the healthy type's IC constraint does not bind since, by Eq. (3), $I_1(\theta_h,\theta_h) \geq r(\theta_h)$. Consequently, both $\bar{t}_1(\theta_h)$ and $t_1(\theta_h)$ are optimally set to $r(\theta_h)$. Reducing the transfers to that level increases the principal's expected welfare and weakens the remaining IC constraint by reducing $I(\theta_h,\theta_s)$. Incorporating these results, the sick type's IC constraint reduces to:

$$G_1(\theta_s)\max\{\bar{t}_1(\theta_s), r(\theta_s)\} + [1 - G_1(\theta_s)]\, r(\theta_s) \geq G_1(\theta_h)\, r(\theta_s) + [1 - G_1(\theta_h)]\, r(\theta_h). \tag{7}$$

This constraint must be binding at the optimum, else $\bar{t}_1(\theta_s)$ could be slightly reduced, improving the principal's welfare without violating any other constraint. There are two relevant cases with this constraint. First, if healthy declarations are always audited, it is optimal to set $\bar{t}(\theta_s) = 0$ since reductions in this transfer improve the principal's welfare without violating IC. If $\bar{t}(\theta_s) = \underline{t}_1(\theta_s) = t_1(\theta_s) = 0$, however, there is no reason to undertake costly audits of sick declarations, thus contradicting our original supposition. Consider the second case in which healthy declarations are audited with some probability less than one. In that case IC requires $\bar{t}_1(\theta_s) > r(\theta_s)$. Then, the sick type's IC constraint reduces to:

$$G_1(\theta_s)\,\bar{t}_1(\theta_s) + [1 - G_1(\theta_s)]\,r(\theta_s) = G_1(\theta_h)\,r(\theta_s) + [1 - G_1(\theta_h)]\,r(\theta_h)\,.$$

The principal's problem can then be re-expressed fully incorporating all constraints as:

$$\begin{aligned}&\max_{G_1(\theta_h),G_1(\theta_s)} [1-\rho]\{P(\theta_h) - r(\theta_h) - G_1(\theta_h)\,a\} + \\ &\rho\{P(\theta_s) - [G_1(\theta_h)\,r(\theta_s) + [1 - G_1(\theta_h)]\,r(\theta_h) + G_1(\theta_s)\,[r(\theta_s) + a] - r(\theta_s)]\}\end{aligned} \tag{8}$$

Since this objective is linearly decreasing in $G_1(\theta_s)$, at optimum $G_1(\theta_s) = 0$. This result also contradicts the original supposition that $G_1(\theta_s) > 0$. Consequently, it is never optimal to audit sick declarations.

Having established that optimally $G_1(\theta_s) = 0$, note from Eqs. (3) and (4) that $I_1(\theta_s, \theta_s) = I_1(\theta_h, \theta_s)$ at the optimum. Otherwise, $t_1(\theta_s)$ could be reduced, improving the principal's welfare without violating any other constraints. This IC constraint can thus be simplified to:

$$\max\{t_1(\theta_s)\,, r(\theta_s)\} = G_1(\theta_h)\,r(\theta_s) + [1 - G_1(\theta_h)]\max\{t_1(\theta_h)\,, r(\theta_s)\}\,. \tag{9}$$

In turn, this result implies that the healthy type's IC constraint can safely be ignored as long as its voluntary participation constraint is satisfied. In other words, $I_1(\theta_h, \theta_h) \geq r(\theta_h) \Rightarrow I_1(\theta_h, \theta_h) \geq I_1(\theta_s, \theta_h)$. To see this, note that using Eqs. (4) and (9),

$$I_1(\theta_s, \theta_h) = \max\left\{G_1(\theta_h)\,r(\theta_s) + \left[1 - G_1\left(\tilde{\theta}\right)\right]\max\{t_1(\theta_h)\,, r(\theta_s)\}\,, r(\theta_h)\right\} \tag{10}$$

Consequently, the principal's welfare can be improved by reducing both the audited and unaudited healthy payments to: $\bar{t}_1(\theta_h) = t_1(\theta_h) = r(\theta_h)$. Using this result, constraint (9) becomes

$$\max\{t_1(\theta_s), r(\theta_s)\} = G_1(\theta_h) r(\theta_s) + [1 - G_1(\theta_h)] r(\theta_h). \tag{11}$$

Thus, if $G_1(\theta_h) = 1$, $t_1(\theta_s)$ is optimally reduced to zero. Otherwise, $t_1(\theta_s) = G_1(\theta_h) r(\theta_s) + [1 - G_1(\theta_h)] r(\theta_h)$. For $G_1(\theta_h) = 1$, the principal's welfare is simply:

$$[1 - \rho][P(\theta_h) - r(\theta_h) - a]. \tag{12}$$

In general, $G_1(\theta_h)$ is chosen to maximize:

$$\begin{aligned} &[1 - \rho]\{P(\theta_h) - r(\theta_h) - G_1(\theta_h) a\} + \\ &+ \rho\{P(\theta_s) - [G_1(\theta_h) r(\theta_s) + [1 - G_1(\theta_h)] r(\theta_h)]\} \end{aligned} \tag{13}$$

Since (13) is linear in $G_1(\theta_h)$, if it is optimal to audit at all it is optimal to audit with probability equal to unity. Auditing is optimal if (13) is increasing in $G_1(\theta_h)$. That is to say, if:

$$\rho[r(\theta_h) - r(\theta_s)] > [1 - \rho] a \tag{14}$$

Intuitively, the left side of this expression represents the expected marginal benefit of auditing. Without an audit, sick types would receive an payment equal to the healthy types' reservation income. With an audit, sick types would receive their own reservation income. The benefit of the audit is the difference in payments to sick types. The right hand side of the expression is simply the expected marginal cost of increasing audits for healthy types.

If, however,

$$\rho[r(\theta_h) - r(\theta_s)] < [1 - \rho] a \tag{15}$$

then the optimal probability of auditing a healthy type approaches zero. For $G_1(\theta_h) = 0$, Eq. (11) implies $t_1(\theta_h) = t_1(\theta_s)$. Intuitively, with no auditing the principal has no means of discriminating between sick and healthy types, so all must receive the same payment in order to reveal the truth. The principal's expected welfare from purchasing a lamb is then:

$$\rho P(\theta_s) + [1 - \rho] P(\theta_h) - r(\theta_h). \tag{16}$$

Here we have the similarity with Akerlof (1970). If the proportion of sick sheep in the population is high enough, the expected value to the principal of purchasing a sheep will be negative. In such a case, the principal will not purchase any sheep, even though with perfect information he could profitably buy them from healthy types.

If the interaction between the agents and the principals were characterized by perfect persistence (i.e., agents with sick sheep in one period have sick sheep in the next period) and credible commitment, one would expect two possible contracts, both exhibiting deterministic auditing and truthful revelation of sheep health.[6] In the first outcome, agents would truthfully indicate the type of their sheep, agents with healthy sheep would receive a higher payment, and all healthy declarations would be audited. In the second, agents would truthfully indicate the type of their sheep, all agents would receive the same payment, and no declaration would be audited.

Actual market conditions do not resemble either of these outcomes. Audits of sheep health are never undertaken (which is not surprising due to the high cost of audits). All sheep receive the same payment. Agents do not, however, voluntarily indicate the type of their sheep. As a consequence, sick sheep occasionally are purchased resulting in human illness, risks of boycotts, and generally suboptimal outcomes for everyone. Cases occur, such as in the Akerlof (1970) model, where the probability of illness is sufficiently hight that the market breaks down altogether.

One reason agents may be reluctant to indicate the type of their sheep is that they are afraid that the principal will use this information against them in the future. If the purchaser cannot credibly commit to not use this information then the solutions described above are not necessarily equilibria. The next section characterizes the optimal mechanism without the possibility of commitment.

[6] This is in contrast to Townsend (1979), where the assumption risk averse preferences led him to conjecture that random audit strategies would always dominate deterministic ones.

2.2 No Commitment

In the case examined in this section, the principal is unable to credibly commit to following a second period strategy that is against his interest at the time. Unlike the previous case, the agent thus believes that by revealing his type in the first period, the principal will take advantage of this information in the second period by not purchasing from a sick type and paying a healthy type only his reservation income. Thus, the second period equilibrium payoffs for a truthful direct revelation mechanism are simply $t_2(\theta_s, \theta_s) = 0$ and $t_2(\theta_h, \theta_h) = r(\theta_h)$. Moreover, since the principal knows agent types, there is no reason to audit in the second period, so $G_2(\theta, \theta) = 0$ for all current and previous declarations of type. Similar to the full commitment case, since transfers are costly to principal and the penalty only appears on the right-hand-side of IC constraints, the optimal penalty transfer will be as small as possible, i.e., $\underline{t}_1(\theta_s) = \underline{t}_1(\theta_h) = 0$.

When confronted with a menu of contracts, in the first period an agent chooses the one that maximizes the present value of his expected income, knowing how the principal will react to his declaration in the next period. For this model, let $I_2\left(\tilde{\theta}, \theta\right)$ denote the present value of expected income for an agent of type θ choosing the contract intended for an agent of type $\tilde{\theta}$. With voluntary ex post participation we have:

$$I_2(\theta_s, \theta_s) = G_1(\theta_s) \max\{\bar{t}_1(\theta_s), r(\theta_s)\} + [1 - G_1(\theta_s)] \max\{t_1(\theta_s), r(\theta_s)\} + \delta r(\theta_s) \quad (17)$$

$$I_2(\theta_h, \theta_s) = G_1(\theta_h)[1 + \delta] r(\theta_s) + [1 - G_1(\theta_h)] [\max\{t_1(\theta_h), r(\theta_s)\} + \delta r(\theta_h)] \quad (18)$$

$$I_2(\theta_h, \theta_h) = G_1(\theta_h) \max\{\bar{t}_1(\theta_h), r(\theta_h)\} + [1 - G_1(\theta_h)] \max\{t_1(\theta_h), r(\theta_h)\} + \delta r(\theta_h) \quad (19)$$

$$I_2(\theta_s, \theta_h) = G_1(\theta_s) r(\theta_h) + [1 - G_1(\theta_s)] \max\{t_1(\theta_s), r(\theta_h)\} + \delta r(\theta_h) \quad (20)$$

For the direct revelation mechanism to be truthful, each agent must voluntarily to choose his own contract. That is to say, the following IC constraint must be satisfied for each type: $I_2(\theta, \theta) \geq I_2\left(\tilde{\theta}, \theta\right)$.

The no-commitment equilibria differ from the full commitment equilibria in one key aspect. Like the previous case, full auditing of healthy declarations may be optimal if audit costs are sufficiently low. Unlike the previous case, however, the complete absence of auditing is not an equilibrium. Rather, random audits of healthy declarations can be optimal. In the likely case that random audits are optimal, the payment scheme differs markedly from the full commitment equilibrium. Namely, the first period transfer to unaudited healthy declarations can be strictly *lower* than for sick declarations. Audited healthy declarations, however, receive a bonus payment. It is the expected value of this bonus that gives healthy types an incentive to tell the truth.

First, note that similar to the full commitment setting, it is never optimal to audit sick declarations with strictly positive probability. Since the proof is essentially the same as in the previous section, it will not be repeated here. Given $G_1(\theta_s) = 0$, the sick type's IC constraint must be binding since otherwise $t_1(\theta_s)$ could be reduced and weaken the healthy type's IC constraint. The sick type's IC constraint simplifies to:

$$I_2(\theta_s, \theta_s) = I_2(\theta_h, \theta_s) \Rightarrow$$
$$\max\{t_1(\theta_s), r(\theta_s)\} + \delta r(\theta_s) = G_1(\theta_h)[1+\delta] r(\theta_s) + [1 - G_1(\theta_h)][\max\{t_1(\theta_h), r(\theta_s)\} + \delta r(\theta_h)]. \quad (21)$$

Since for $G_1(\theta_h) < 1$ the right-hand side of Eq. (21) is greater than $[1+\delta] r(\theta_s)$, for any healthy audit probability less than one it must be the case that $t_1(\theta_s) > r(\theta_s)$. Moreover, the principal can reduce $t_1(\theta_h)$ to $r(\theta_h)$ without adversely affecting her welfare. Such a move reduces the size of $t_1(\theta_s)$ necessary to satisfy Eq. (21). If accompanied by a corresponding increase in $\bar{t}_1(\theta_h)$ it does not increase the expected cost of satisfying the healthy type's IC constraint. Thus, for $G_1(\theta_h) < 1$, Eq. (21) can be simplified further to:

$$t_1(\theta_s) = G_1(\theta_h)[1+\delta] r(\theta_s) + [1 - G_1(\theta_h)] r(\theta_h) + \delta[r(\theta_h) - r(\theta_s)]. \quad (22)$$

The healthy type's IC constraint then simplifies to:

$$I_2(\theta_h, \theta_h) \geq I_2(\theta_s, \theta_h) \Rightarrow$$
$$G_1(\theta_h) \max\{\bar{t}_1(\theta_h), r(\theta_h)\} + [1 - G_1(\theta_h)] r(\theta_h) \geq \max\{t_1(\theta_s), r(\theta_h)\}. \quad (23)$$

Note that this constraint is binding. For any $\bar{t}_1(\theta_h) > r(\theta_h)$ for which the left-hand side of (23) is strictly greater than the right-hand side, the principal can improve her expected welfare by holding all else constant and reducing $\bar{t}_1(\theta_h)$ without affecting other constraints. It is never optimal, however, for the principal to reduce $\bar{t}_1(\theta_h)$ below $r(\theta_h)$. Doing so would not weaken any IC constraints. In addition, it would cause audited healthy types to refuse to participate. The principal would then forfeit income equal to $P(\theta_h) - r(\theta_h)$ for those agents. The healthy types' IC constraint can thus be simplified to:

$$G_1(\theta_h) \bar{t}_1(\theta_h) + [1 - G_1(\theta_h)] r(\theta_h) = \max\{t_1(\theta_s), r(\theta_h)\}. \quad (24)$$

By comparing contraints (22) and (24) we see that any contract that entails no auditing cannot be incentive compatible. For by (22), $G_1(\theta_h) = 0 \Rightarrow t_1(\theta_s) > r(\theta_h)$. Yet by (24), $G_1(\theta_h) = 0 \Rightarrow t_1(\theta_s) \leq r(\theta_h)$.

From (22) and (24) we can also see that if $G_1(\theta_h)$ is changed by a small amount $dG_1(\theta_h)$, then incentive compatibility is preserved by changing $\bar{t}_1(\theta_h)$ in the opposite direction by an amount $d\bar{t}(\theta_h)$. If $t_1(\theta_s) > r(\theta_h)$, then

$$d\bar{t}(\theta_h) = \frac{r(\theta_s) - \bar{t}(\theta_h)}{G_1(\theta_h)} dG_1(\theta_h). \tag{25}$$

If $t_1(\theta_s) \leq r(\theta_h)$, then

$$d\bar{t}(\theta_h) = -\frac{\bar{t}(\theta_h)}{G_1(\theta_h)} dG_1(\theta_h). \tag{26}$$

In other words, for the principal there is a tradeoff between the size of the reward for truth-telling and the probability of audit. For any incentive compatible combination of $\bar{t}_1(\theta_h)$ and $G_1(\theta_h)$, another combination with higher (lower) $\bar{t}_1(\theta_h)$ and lower (higher) $G_1(\theta_h)$ will also be incentive compatible.

Using the results obtained thus far in this section, the principal's welfare function can be simplified from Eq. (2) to:

$$\begin{aligned} W_2 = {} & [1-\rho]\{[1+\delta] P(\theta_h) \\ & - [G_1(\theta_h)[\bar{t}_1(\theta_h) + a] + [1 - G_1(\theta_h) + \delta] r(\theta_h)]\} \\ & + \rho\{P(\theta_s) - [G_1(\theta_h) r(\theta_s) + [1 - G_1(\theta_h)] r(\theta_h) + \delta[r(\theta_h) - r(\theta_s)]]\} \end{aligned} \tag{27}$$

Differentiation of (27) shows that from any initial combination in which IC constraints are satisfied, the gain to the principal's welfare from a marginal increase in $G_1(\theta_h)$ and decrease in $\bar{t}_1(\theta_h)$ that satisfies IC is:

$$[r(\theta_h) - r(\theta_s)] - [1-\rho] a \tag{28}$$

The first bracketed term is the benefit of the information obtained from auditing in terms of reduced transfers to both types of agents. Notice that this benefit is higher than in the full commitment case. The reason is that incentive compatibility constraints bind for both types without commitment as opposed to only binding for the low type with full commitment. The second term is the expected cost of a marginal increase in audit probability. If expression (28) is positive, then it is optimal to increase the probability of audit as high as possible, i.e., to unity. If the expression is negative, it is always optimal to reduce $G_1(\theta_h)$. As discussed earlier, however, the probability of audit cannot be reduced to zero and still satisfy IC constraints.

If there is an upper bound $t^{\max} > r(\theta_h)$ on the credible payment for being "caught" telling the truth, then $G_1(\theta_h)$ can only be reduced until the upper bound on transfers is reached. In this case, the optimal value of $G_1(\theta_h)$ is obtained by rearranging Eq. (24) to solve

$$G_1(\theta_h) = \frac{\delta[r(\theta_h) - r(\theta_s)]}{t^{\max} - [1+\delta] r(\theta_s)}. \tag{29}$$

To summarize, there are two classes of equilibra to the no-commitment game. Neither is characterized by auditing sick declarations. If the cost of auditing is sufficiently low, the principal always audits healthy declarations. In this case, each type of agent receives his reservation income. The other class of equilibria is characterized by a random auditing of healthy declarations. In this case, healthy agents receive their reservation income if not audited. Audited, healthy agents receive an bonus payment. The higher this additional payment, the lower is the necessary probability of audit. Sick agents receive payments strictly greater than their reservation income in the first period. In the second period, all types receive their reservation income.[7] If the audit cost is sufficiently high, and the maximum bonus is sufficiently low, the market may still collapse since the principal's expected profit would be negative.

One reason agents may be unwilling to disclose the type of their sheep is the fear that the principal will later use this information against them. The analysis in this section shows that if this is the case, it may be optimal to encourage truth-telling by offering a premium for sick sheep, and a bonus for audited healthy sheep.

3 Conclusion

In addressing the problem of the normalization of the livestock trade disrupted by Rift Valley Fever, we have presented a core model expressing the basic incentive issues resulting from private information, producer reputation, and buyer credibility. The analysis implies that a cooperative outcome might be achieved if the principal rewards the agent for revealing private health information. The feasibility of the outcome depends on the probability of sickness, the cost of monitoring, and the values of healthy and sick sheep for the players.

[7] These results are similar to those obtained by Border and Sobel (1987) for multiple types in the single period game.

In the equilibrium contract identified, the principal gains by setting incentives so that being audited is like winning a lottery–if the agent is truthful. The agent wins by being "lucky enough" to be found to be telling the truth. This contract eliminates situations in which it is worthwhile for the agent to declare that a sick sheep is healthy. It also eliminates situations in which it is worthwhile for the agent to attempt to benefit from rewards by declaring that healthy sheep are sick.

Reputation of the agent and credibility of the principal are central to the problem. If the principal can credibly commit to ignoring information revealed by an agent in one period when taking actions in following periods, then a broader range of cooperative outcomes are available. However, it is not clear that the principal could commit to ignoring information, that once obtained, could be used to increase profits.

In addition, in the equilibrium identified, if the principal performs expensive auditing and provides large rewards to those reporting truthfully, then it is in the interest of the agent to report truthfully. However, even though the sellers are reporting truthfully, the principal is expending a great deal of resources for testing when she knows what the outcome of the tests will be. Therefore, the principal must be able to commit to performing the expensive auditing even when she knows that everyone is reporting truthfully. This may also be a commitment that is difficult for the principal to make credibly. In further research using a model similar to that of Khalil (1997) we will examine in greater detail the problem of the principal being unable to commit to an auditing strategy that is not time consistent.

We have performed the first of many modeling steps necessary to answer how, and if the livestock market could be normalized. Continuing our work, we will attempt to find the conditions in which people would not block a forecasting and monitoring system. We will study if third party participation is not only necessary to implement technical aspects of the system, but also to ensure credibility in contracts. It may be that a system for a credible commitment to monitoring, rewards for truthful reporting, and the protection of private information must be designed into the forecasting an monitoring system during negotiations concerning its implementation.

We plan to examine a richer model in which an agent type is imperfectly persistent over time. Rather than having a fixed type, sheep health follows a Markov process in which its future probability distribution depends on its current state. Such analysis would allow us to examine the impact of climate forecasting tools that will improve the precision with which the principal develops her beliefs regarding agent type. Through this additional analysis, we hope to provide insight to the following question: Although forecast and monitoring systems are necessary for a desirable class of cooperative equilibria, can this same information prevent cooperation if its quality is too high?

Other extensions include allowing for possible geographic correlation among agent types. In this model it may be worthwhile for the principal to make the contracts contingent not just on an individual's declaration, but also on his neighbors' declarations. Such a model would allow for a richer strategic interaction among the agents in terms of determining the equilibrium strategy. This work could motivate incentives for group effects, potentially leading to contracts that utilize group effects for improved outcomes.

As a further extension, we will examine the impact of RVF forecasts on the market. An interesting feature of the problem is that preventative measures, such as immunization are possible but costly. This extension would include moral hazard elements similar to Mookherjee and Png (1989). Further complications arise, however, since immunization must take place weeks before sheep are exposed to the disease. Nomadic pastoralists have very little access to the Internet and communication technologies, so, in some cases only the buyers may have access to forecast information. If there is differential access to forecast products, this may lead to equity and management issues. We will examine if the buyers might communicate this information through the contract to provide producers in afflicted regions with signals about when to immunize. We also hope to identify and understand the ways in which forecasts might end up being used by buyers against the sellers when sellers do not have the same level of access as buyers.

In addition to being an academic contribution, results from this work are of policy relevance. African producer interests are currently engaged in contentious negotiations with Middle Eastern importers in order to establish a forecast system and prevent boycotts. Through our work, we hope to inform the debate concerning the potential for the success of monitoring and disease forecasting systems and improving the chances of successful negotiation outcomes by illustrating the potential for Pareto improvements and informing third parties of additional roles they may need to play.

References

Akerlof, G. A., 1970. The market for "lemons": Quality uncertainty and the market mechanism. Quarterly Journal of Economics 84 (3), 488–500.

Baron, D., Besanko, D., 1984. Regulation and information in a continuing relationship. Information Economics and Policy 1, 447–470.

Border, K. C., Sobel, J., 1987. Samurai accountant: A theory of auditing and plunder. Review of Economic Studies 54, 525–540.

Consultative Meeting of Experts on Rift Valley Fever, 2001. Reducing the Risk of Rift Valley Fever Virus Transmission. Report, United Nations FAO and UNDP.

Dilley, M., 2003. Rift Valley Fever (RVF) Risk Model: Design Process. Report, Columbia University IRI, AU/IBAR, USAID.

Indeje, M., Ward, N., Ogallo, L. J., Davies, G., Dilley, M., Anyamba, A., 2004. Predictability of the Normalized Difference Vegetation Index (NDVI) in Kenya, An Indicator of Rift Valley Fever Outbreaks in the Greater Horn of Africa. Working paper, Columbia University IRI.

Khalil, F., 1997. Auditing without commitment. Rand Journal of Economics 28 (4), 629–640, times Cited: 13 Article English Cited References Count: 20 Yr264.

Mookherjee, D., Png, I., 1989. Optimal auditing, insurance, and redistribution. Quarterly Journal of Economics 104 (2), 399–415, times Cited: 82 Article English Cited References Count: 12 U6400.

Osgood, D., 2005. Snowblind: The importance of climate information for recreational real estate. Working paper, Columbia University IRI.

Pfaff, A., Broad, K., Glantz, M., Feb 1999. Who benefits from climate forecasts? Nature 397, 645–646.

Townsend, R. M., 1979. Optimal-contracts and competitive markets with costly state verification. Journal of Economic Theory 21 (2), 265–293.

Townsend, R. M., 1988. Information constrained insurance - the revelation principle extended. Journal of Monetary Economics 21 (2-3), 411–450.

Wang, C., 2005. Dynamic costly state verification. Economic Theory 25 (4), 887–916.